Stockholm School of Economics
4210 - Thesis within Finance
February 2007

On the Informative Content of Technical Indicators

- a Systematic Approach to Pattern Recognition

Joel Mossberg[a] Johan Undin[b]

Abstract

Technical analysis is a rather subjective graphical method of recording the history of trading in a certain stock to deduct the probable future trend in its return. By applying the automatic computer algorithm proposed by Lo, Mamaysky and Wang (2000) on Swedish stock market data we separate ourselves from the subjectivity of ordinary pattern recognition. In this way we are able to systematically determine the informative content of ten commonly used technical indicators over the time-period of 1982 - 2006. At the same time we also perform a large scale out-of-sample test of the ability of the Lo, Mamaysky and Wang (2000) algorithm to detect technical patterns. We conclude that the Lo, Mamaysky and Wang (2000) model performs well, also on Swedish data, and that many of the technical patterns indeed contain incremental information which could be useful for an investment strategy.

[a] 20248@student.hhs.se
[b] 20154@student.hhs.se

Contents

1 Introduction

1.1 Technical Analysis

"The art of technical analysis, for it is an art, is to identify stock market movement trends at an early stage and to open an investment posture when a reversal of that trend is indicated." (Pring, 1980)

Technical analysis is the method of forecasting future stock prices from geometric shapes in historical stock price movements and is also known as charting. Technical analysis has played a role in investment strategies ever since Edwards and Magee wrote their influential book "Technical Analysis of Stock Trends" in 1948, now considered the corner stone of pattern recognition analysis. However, technical analysis has not been given the same general acceptance as e.g. fundamental analysis among the broader public of investors, both academics and practitioners. One reason for this is the rather subjective side of technical analysis, where pattern recognition often is in the eyes of the beholder. Technical analysis has long been considered the black sheep of investment strategies and it has been given the not very flattering nick names of voodoo finance and financial alchemy.

One of the most plausible reasons for the contempt for technical analysis from the academic critics lie in the fact that technical analysis is based on visual judgements whereas quantitative finance is mainly algebraic and numerical. It follows that quantitative finance is less abstract and can be practiced with the help of computers, numerical algorithms, mathematics and statistics.

In order to close the gap between technical analysis and quantitative finance attempts have been made to develop systematic and scientific approaches for the theory of technical analysis. In this paper we are using such a method in order to conduct a computerized analysis when performing technical analysis on Swedish stocks.

1.2 Research Objective

Our aim in this paper is to apply the same systematic and automatic approach to the practise of technical analysis using nonparametric kernel regression as do Lo, Mamaysky and Wang (2000), but we intend to test the algorithm on Swedish stock data.

By applying the same algorithm as Lo, Mamaysky and Wang (2000), our study can be seen as a large scale out-of-sample test of their methodology. As technical analysis is somewhat built around data mining, there are always reasons to believe that a certain algorithm performs very well on one sample, but lacks the general capability to repeat the results for other data sets. If the Lo, Mamaysky and Wang (2000) algorithm should be seen as general and if there is some substance in technical indicators, we should expect fairly similar results on our data set. At the same time as testing the Lo, Mamaysky and Wang (2000) model we test the informative content of technical indicators in Swedish stock data. It is important to notice that the efficiency of the Lo, Mamaysky and Wang (2000) model and the informative content in technical indicators are tested simultaneously and only if both works we will get statistically significant results. Finally, we intend to wrap up with a short discussion of the usefulness of technical indicators as a trading tool and its implications for the efficient market hypothesis.

As far as we are concerned no empirical study of this type has been conducted on Swedish data, making our study the first of its kind and therefore contributing to the research community. It is also worth noticing that this paper must not be seen as a support for the trading strategies of technical analysis. It is rather the opposite way. We limit ourselves to apply the Lo, Mamaysky and Wang (2000) algorithm on Swedish data, leaving for the readers to make up their own mind about technical analysis as an investment tool.

1.3 Thesis Outline

The outline of this thesis can be summarized as follows. In the first section we give a short introduction to the subject of the thesis and our research objectives. In the second section we provide an overview of the previous research done on the subject. We discuss the efficient market hypothesis, general support for technical analysis and the Lo, Mamaysky and Wang (2000) paper. In the third section we map out the theoretical framework and the major analytical tools related to the study. We describe the foundations and the important patterns of technical analysis. We provide an introduction to smoothing estimators and a thorough explanation of the specific smoothing estimator used in this study, namely the kernel regression. We also introduce the statistical distribution tests that we use. In the forth section we provide a detailed description of our methodology. We describe the data we have used and our algorithm for automating technical analysis. The analysis of our results is provided in section five. Finally, we summarize in section six.

2 Previous Research

In order to describe the previous research conducted in the area of technical analysis, we aim to take a funnel approach. Since technical analysis by definition is nonsense if the market indeed is efficient, we start with a discussion of the efficient market theory and focus on evidence discarding the hypothesis. We will then narrow the discussion to more direct support for technical analysis, trying to give a brief overview of the research conducted on the subject so far. Finally, we give an overview of one specific article of great importance to our paper - Foundations of Technical Analysis: Computational Algorithms, Statistical Inference and Empirical Implementations by Andrew W. Lo, Harry Mamaysky and Jiang Wang (2000).

2.1 Evidence on the Efficient Market Hypothesis

The search for predictability in asset returns has been on investors' and academics' minds since the birth of organized financial markets. It follows that a great deal of research has been conducted on examining the efficiency of stock market price formations. In recent years, doubts have been raised about the efficient market hypothesis, something that would give indirect support to the possibility of technical analysis.

A number of papers suggest that stock returns are not fully explained by common risk measures. For example, a significant link between expected return and fundamental variables such as price-earnings ratio, market-to-book ratio, and size has been found (see, for example, Basu (1977) and Fama and French (1992)). Another group of papers have brought in light systematic patterns in stock returns related to various calendar events such as the weekend effect, the turn-of-the-month effect and the January effect (see, for example, Schwert (2003) and Lakonishok and Schmidt (1988)).

Another direction of research directly related to the theory explained above provides evidence of predictability of equity returns from past returns. This is the so called weak form of market efficiency (see, for example, Fama (1991)). The fact that past prices can be used to predict future returns is the sole base for technical analysis and a fact all technical analysts take for granted. In fact, when studying the academic work on this subject, the expression technical analysis is often used as a wider expression for past prices containing information of future returns. Lo and MacKinlay (1988) strongly rejects the random walk model and suggests that the returns of their created portfolios are auto-correlated in a paper that tests the random walk hypothesis for weekly stock market returns. Treynor and Ferguson (1984) show that past prices, when combined with other valuable information, can indeed be helpful in achieving unusual profit. Brown and Jennings (1989) use a two-period dynamic model of equilibrium to demonstrate that rational investors use historical prices in forming their demands and to illustrate the sensitivity of the value of technical analysis to changes in the values of exogenous parameters.

Many papers have documented that average stock returns are related to past performance. Jegadeesh and Titman (1993) documents that strategies which buy stocks that have performed well in the past and sell stocks that have performed poorly in the past generate significant positive returns over a three to twelve month trading window. DeBond and Thaler (1985, 1987) states that stock prices overreact to information suggesting that buying past losers and selling past winners achieve excess returns. In line with previous papers, De Bondt and Thaler (1985) also show that over a three to five year holding period stocks that performed poorly over the previous three to five years achieve higher returns than stocks that performed well over the same period. Chan, Jegadeesh and Lakonishok (1996) examines whether the predictability of future returns from past returns is because of the market's underreaction to information, in particular to past earnings news. Their results point in the direction that we have a market responding only gradually to new information. Rouwenhorst (1998) shows that between 1980 and 1995 an internationally diversified portfolio of past medium-term winners outperforms a portfolio of medium-term losers after correcting for risk by more than 1 percent per month.

Blume, Easley and O'Hara (1994) investigate the informational role of volume, something considered being of great importance to technical analysts. They show that volume provides information on information quality that cannot be deducted from the price statistics. They also show how volume, information precision, and price movements relate, and demonstrate how sequences of volume and prices can be informative.

2.2 Support for Technical Analysis

Technical analysis has experienced surging support both among practitioners and in the academic world (Lo and MacKinlay, 1999). More or less direct support for technical analysis has been given in a number of studies. One of the first illustrations of technical analysis is the discussion of Dow Theory in Rhea (1932). In far more recent studies, Pruitt and White (1988) try to directly determine the profitability performance of a technical trading system including price, volume, and relative strength indicators on individual stock issues. The study shows that the trading system has the ability to beat a simple buy-and-hold strategy over a significant period of time and therefore generates support for technical analysis. Neftci (1991) investigates statistical properties of technical analysis in order to determine if there is any objective foundation for the attractiveness of technical pattern recognition. The paper examines whether formal algorithms for buy and sell signals similar to those given by technical analysts can be made and whether the rules of technical analysis are useful in prediction in excess of the forecasts generated by the Wiener-Kolmogorov prediction theory. The article shows that most patterns used by technical analysts need to be characterized by appropriate sequences of local minima and/or maxima and if defined correctly, technical analysis can be useful over and above the Wiener-Kolmogorov prediction theory. Brock, Lakonishok and LeBaron (1992) test two of the simplest and most popular trading rules, moving average and trading range break (resistance and support levels). The result shows strong support for the technical strategies when compared with four popular null models: the random walk, the AR(1), GARCH-M, and the Exponential GARCH. Dittmar, Neely and Weller (1997) are using genetic programming techniques to find technical trading rules and find strong evidence of economically significant out-of-sample excess returns.

2.3 The Lo, Mamaysky and Wang (2000) Paper

The common denominator in the papers discussed in the previous section are that they all use fairly simple patterns, e.g. crossing trend lines, and compare them to a theoretical development of stock prices. The one exception is Neftci (1991) who showed that, in principle, all technical analysis patterns can be formally defined using particular sequences of local minima of maxima. However, he does not show whether it gives excess return, only that it has predictive power.

Following the reasoning of Neftci (1991), there is one special article that opened our eyes since it is the only one in our knowledge that with the help of computer algorithms tries to replicate the work of a trained technical analyst. Lo, Mamaysky and Wang (2000) propose a systematic and automatic approach to technical pattern recognition and apply the method on US stock data to evaluate the effectiveness of technical analysis. In order of doing so, they first need to extract nonlinear patterns from noisy data, i.e. the stock price development. They use a class of statistical estimators called smoothing estimators for this task. More specifically, they use a specific type called kernel regression. Using this regression they apply their algorithm for automating technical analysis on the daily returns of several hundred US stocks from 1962 to 1996. They find that over the 31-year sample period, several technical indicators do provide incremental information and are in fact suggesting that technical analysis may have some practical value.

3 Theoretical Framework

3.1 On Technical Analysis

In its application to the stock market, the term technical has come to have a very special meaning, quite different from what we normally interpret into the term. It refers to the study of the market itself as had it its own life. Technical analysis is "the science of recording, usually in graphic form, the actual history of trading (price changes, volume of transactions, etc.) in a certain stock and then deducting from that pictured history the probable future trend" (Edwards and Magee, 1997).

The general goal of technical analysis is to identify regularities in the time-series of prices by extracting nonlinear patterns from noisy data. It follows, that charts are the working tools of the technical analyst. Volume, or the trading activity in a stock, also plays an important role for the technical analyst (see, for example, Blume, Easley and O'Hara (1994)).

In this paper, and following Lo, Mamaysky and Wang (2000), we have focused on five pairs of technical patterns that are considered being among the most popular and famous patterns of traditional technical analysis. These are head-and-shoulders and inverted head-and-shoulders, broadening tops and bottoms, triangle tops and bottoms, rectangle tops and bottoms, and double tops and bottoms. Of course, there are many other technical indicators that could have been used, some of which might be easier to find with the help of computers. Famous once include moving averages, flags, gaps and trend lines (see, for example, Edwards and Magee (1997)). Since we, following Lo, Mamaysky and Wang (2000), use an advanced computer method to determine non-linear patterns in stock price charts, our focus lies on those patterns that are most difficult to quantify analytically and where the proposed method provides the most value added.

One must keep in mind that the definitions of the various patterns are a bit random and subject to the opinion of different authors. This of course makes it harder to conduct a general study on technical indicators. Nevertheless, we have chosen to use the definitions proposed by Edwards and Magee (1997) and Nilsson and Torssell (2000) combined with Lo, Mamaysky and Wang (2000). For graphical representations of the patterns, please see page 11.

3.1.1 Head-and-Shoulders

The head-and-shoulders pattern is not only the most famous, but also one of the more common and, by all odds, considered the most reliable of the major patterns. It can appear in two ways, as normal head-and-shoulders or as inverse head-and-shoulders.

The normal head-and-shoulders pattern consists of four parts: the two shoulders, the head and the break-out. It starts with a strong upward trend during which the trading volume becomes very heavy, followed by a minor recession on which trading volume decreases. This is the left shoulder. The next section starts with another high-volume rally which reaches a higher level than the top of the left shoulder, and then another downturn on less volume which take prices down to somewhere near the bottom level of the preceding recession. It can be higher or lower but in any case below the top of the left shoulder. This is the head. Then comes a third increase, but this time during much less volume than that of the first two increases, which fails to reach the height of the head before another decline sets in. This is the right shoulder. Finally, a decrease of the stock price in this third recession down through a line, called the neckline, drawn across the bottoms of the declines on both sides of the head. This decline should close below the neckline by an amount approximately equal to 3 percent of the stock's market price. This is the confirmation or the break-out. The break out of the head-and-shoulders pattern is a signal for selling the stock.

The inverted head-and-shoulders pattern looks the same as the normal one apart from the obvious fact that it is turned upside down. The break out of the inverted head-and-shoulders pattern is a signal for buying the stock.

3.1.2 Broadening Tops and Bottoms

The broadening patterns start with very narrow fluctuations and then widen out between diverging boundary lines. The tops start with a maximum and the bottoms start with a minimum.

The trading activity during a broadening formation usually remains high and irregular throughout its construction. The appearance of this patterns suggest that the market is approaching a dangerous stage indicating that new commitments should not be made and any holdings should be cashed in at the first good opportunity. It is reasonable to assume that the prices, if they break away from the formation, will go down. Thus, by all means, the broadenings are sell signals.

3.1.3 Triangle Tops and Bottoms

Historically, triangles have developed at periods of major trend changes and they are therefore considered as important since these are the periods which are most relevant for an investor to realize. Triangles normally signal a consolidation in the market, terminating an up or down move only temporary and preparing for another strong move in the same direction at a later stage.

The triangle tops are composed by a series of price fluctuations, starting at a maximum, where every new fluctuation is smaller than the last one. This creates a down-slanting line touching the tops of the fluctuations as well as an up-slanting line touching the bottoms. Together, the two lines form a triangle. In the run of this price fluctuation, trading activity shows a decreasing trend. The smaller the fluctuations get, the volume turns into an abnormally low daily turnover. The sign whether to buy or sell comes when the price breaks out of the triangle. This occurs in a notable pick up in volume. If the price increases, it will likely continue doing so and it is therefore a clear buy signal. The opposite goes for a decline. It is very rare that the chart contains any information in which direction the price is going to break out. The investor normally has to wait and see until the action suddenly occurs.

Triangle bottoms are built up in the same way as the tops, with the only difference that they start with a minimum. The buy or sell sign and decision are the same as for the tops.

3.1.4 Rectangle Tops and Bottoms

A rectangle consists of a series of sideways price fluctuations which is called the trading area. It has been given this name since it can be bounded both at the top and at the bottom by horizontal lines. These lines are allowed to slope in either direction if the departure from the horizontal line is trivial. In the same way as for triangles, the rectangle tops starts with a maximum and the bottoms starts with a minimum. The trading volume development within the patterns follows the same rules as for triangles, i.e. the activity decreases as the rectangle lengthens. Also in terms of break outs and indications of directions the same rules as for triangles apply. If the price increases, it will likely continue doing so and is therefore a clear buy signal. The opposite goes for a decline.

3.1.5 Double Tops and Bottoms

The doubles normally occur very rarely and they are difficult to use in the sense that they cannot be detected until prices have gone quite a long way away from them. They can never be told in advance or identified as soon as they occur.

The definition of the doubles is also slightly more involved. The double tops is formed when a stock's price increases to a certain level under heavy trading and then falls back during a decrease in activity. It should then bounce back to approximately the same level as the first top during less heavy trading as last increase. Then, finally, it turns down a second time. The distance between the two tops must not be too small. Lo, Mamaysky and Wang (2000) use a minimum of 23 trading days. The double tops give a signal of selling the stock since the second down turn indicates a consequential decline.

The double bottoms are of course the same pattern turned upside down and it is a signal of buying the stock.

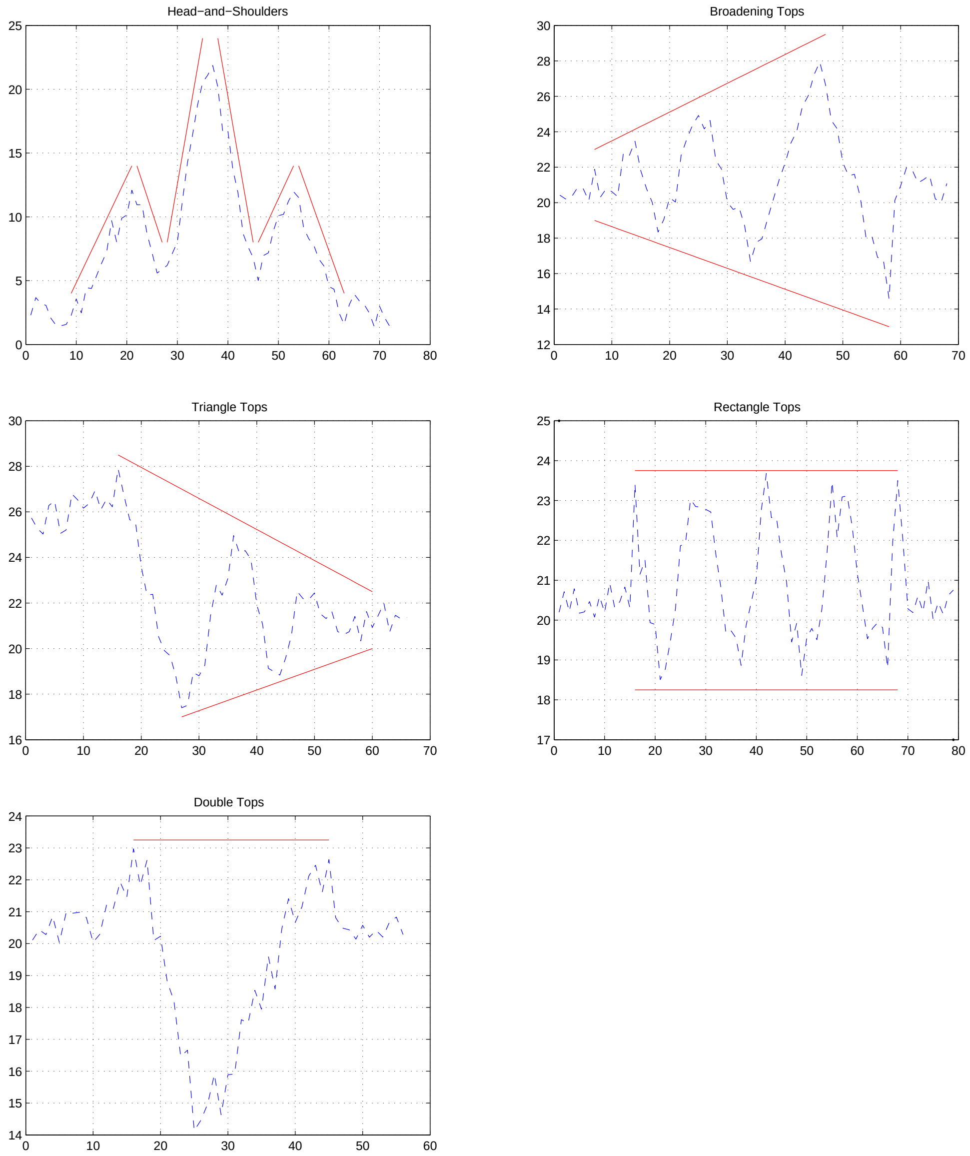

3.2 Smoothing and Kernel Regression

3.2.1 Smoothing Estimators

The natural starting point for any regression, linear or nonlinear, is the regression equation, which is assumed to mirror the behaviour of the underlying variables. For a series of stock prices $\{P_t\}$ the most fundamental equation which captures the nonlinearity is:

$$P_t = m(X_t) + \epsilon_t$$

where $m(X_t)$ is an arbitrary fixed but unknown nonlinear function of a state variable X_t and ϵ_t is white noise. When determining stock prices using time-series data the state variable is usually set equal to time, i.e. $X_t = t$. However, we use the expression $X_t = x_t$ to keep our derivations more in line with Lo, Mamaysky and Wang (2000).

Financial theorists have not yet been able to agree upon a parametrical model for the movement of stock prices, i.e. they have not been able to determine the shape of $m(X_t)$ analytically. The function $m(X_t)$ thus has to be estimated non-parametrically from available data. Lo, Mamaysky and Wang (2000) define pattern recognition as the method of constructing a smooth function $\hat{m}(\cdot)$ to approximate a time-series of prices $\{p_t\}$. The dot indicates that the form of the regression equation does not have to be specified in advance. The fact that the regression equation does not have to be specified, but can be drawn from any data, is an advantage since it does not limit the spectrum of possible patterns which can be found in the data.

One method to estimate the nonlinear function $\hat{m}(\cdot)$ is smoothing, which can be described as a technique to reduce the regression errors by averaging data in some sophisticated way. In a general parametrical regression the $m(\cdot)$ is determined by repeated sampling. By repeating the sampling of $X_t = x_0$ it is possible to determine an estimator of $m(x_0)$ such as:

$$\hat{m}(x_0) = \frac{1}{n}\sum_{i=1}^{n} p_i = \frac{1}{n}\sum_{i=1}^{n}[m(x_0) + \epsilon_t^i] = m(x_0) + \frac{1}{n}\sum_{i=1}^{n}\epsilon_t^i = m(x_0)$$

since $\frac{1}{n}\sum_{i=1}^{n}\epsilon_t^i$ is negligible for large n. Unfortunately, when using time-series we can not allow ourselves to repeat the sampling for a given time t, since only one observation per time-period is available. However, Lo, Mamaysky and Wang (2000) describe a method to avoid this problem. If $m(\cdot)$ is assumed to be sufficiently smooth in a small interval around x_0, then, in a small neighbourhood around x_0, $\hat{m}(x_0)$ will be nearly constant and can be estimated by averaging the P_t's corresponding to those X_t's around x_0. It is obvious that the P_t's closest to x_0 provides more information about

$\hat{m}(x_0)$ than the P_t's further away. Weighting the observations according to some weighting schedule depending on the distance between x_0 and x_t thus improves the estimate. More formally the smoothing estimator of $m(x)$ can be described as:

$$\hat{m}(x) = \frac{1}{T}\sum_{t=1}^{T}\omega_t(x)P_t$$

where the weights $\{w_t(x)\}$ are larger for observations closer to x. The performance of the estimate is to a large extent dependent on the length of the neighbourhood in which $m(\cdot)$ is assumed to be linear and the applied weights. If the neighbourhood is too small and the weights declines too rapidly the regression will be too volatile and too much noise will be captured. On the other hand, if the neighbourhood is too large and the weights too constant, valuable information will be lost. Thus, the weights have to be chosen to balance these two considerations.

3.2.2 Kernel Regression and the Determination of the Estimation Weights

Several methods to determine the regression weights have been proposed in the literature. Härdle (1990) describes a conceptually simple approach called kernel regression estimator. The kernel is defined as a continuous, bounded and symmetrical real function K which integrates to one:

$$K(u) \geq 0, \int K(u)du = 1$$

In order to provide flexibility in terms of the choice of weights so that the above described trade-off between too small and too high weights can be balanced, the kernel is scaled by a factor h so that:

$$K_h(u) = \frac{1}{h}K(u/h), \int K_h(u)du = 1$$

The regression weights are then given by:

$$\omega_{t,h}(x) = K_h(x - X_t)/g_h(x)$$

$$g_h(x) = \frac{1}{T}\sum_{t=1}^{T}K_h(x - X_t)$$

Substituting these weights into the smoothing estimator function yields a kernel estimator $\hat{m}_h(x)$ of $m(x)$:

$$\hat{m}_h(x) = \frac{\sum_{t=1}^{T}K_h(x - X_t)Y_t}{\sum_{t=1}^{T}K_h(x - X_t)}$$

Härdle (1990) shows that $\hat{m}_h(x)$ asymptotically converge to $m(x)$. This converge holds for a wide range of kernels K such as Uniform, Biweight, Triweight, Epanechnikov and Gaussian (Härdle, 1990). The perhaps most commonly used kernel is the Gausian kernel where the kernel K is given by the Gaussian distribution scaled by h (Lo, Mamaysky and Wang, 2000):

$$K_h(x) = \frac{1}{h\sqrt{2\pi}} e^{-\frac{x^2}{2h^2}}$$

3.2.3 Cross-Validation and the Selection of Bandwidth

The scaling parameter is commonly known as the bandwidth of the regression. Too small a bandwidth yields too volatile an estimate while too large a bandwidth conceals valuable information (Fan and Gijbels, 1996). The choice of bandwidth is thus very crucial for the kernel regression.

The bandwidth can of course be set to some constant using a quick look at the data or some rule of thumb. The drawback of this procedure is of course that the choice is very arbitrary and does not necessarily reflect the actual properties of the data. An automatic choice is clearly preferable. Several methods to automatically determine the optimal bandwidth has been suggested. Mittelhammer, Judge and Miller (2000) derive that the bandwidth $h^* = 1.059\sigma n^{-1}$, where σ is the standard deviation of the data, works reasonable for the Gaussian kernel if the data is normally distributed. A more robust method, known as the cross-validation or leave-one-out method, is proposed by Green and Silverman (1994). The method is independent of the distribution of the data and has better finite sample properties (Green and Silverman, 1994).

The cross-validation method is a non-parametric version of the standard method used for parametric regression. The general procedure for determining a regression equation is to train the equation in-sample and then evaluate it out-of-sample. Since a non-parametric regression normally is used on a single data set no new observations are available. Instead, the cross-validation method creates an out-of-sample by omitting one observation at the time and run a regression on the remaining observations (Campbell, Lo and MacKinlay, 1997). Formally, this can be described as minimizing the cross-validation function:

$$CV(h) = \frac{1}{T}\sum_{t}^{T}(P_t - \hat{m}_{h,t})^2$$

where

$$\hat{m}_{h,t} = \frac{1}{T}\sum_{x \neq t}^{T} \omega_{x,h} Y_x$$

The estimator $\hat{m}_{h,t}$ is the kernel estimator applied to the data set with the t-th observation omitted (Green and Silverman, 1994). By selecting the bandwidth h that minimizes the cross-validation function $CV(h)$ the asymptotic mean-squared error is minimized (Baltagi, 2001).

3.3 A Graphical Visualization of Kernel Smoothing Regressions and Automatic Bandwidth Selection

For the reader to gain a deeper understanding of a kernel smoothing algorithm's ability to capture the true trend in noisy data we provide the following clinical example of trend detection from generated data. More specifically, we generate a noisy signal by applying random shocks to a sine-wave and visualize how well the kernel regression can capture this trend. We also demonstrate how different bandwidths affect the regression equation. The equation used to generate the noisy data is:

$$Y = \sin(\frac{16\pi x}{300}) + \frac{x}{50} + 0.3\epsilon \ , \ x = 1, ..., 150$$

where ϵ is a random i.i.d. shock. The upper leftmost figure on page 16 graphs the noisy time-series and the underlying sine-trend. One can clearly see that the data randomly fluctuates around the dashed trend-line. The upper rightmost figure on page 16 provides an example of the kernel method proposed in the theory section together with the cross-validation method for automatic bandwidth detection. We see that the kernel regression function to a large extent mirrors the underlying sine-function, without any parameterisation of the actual shape of the underlying form. This implies that the kernel method finds any underlying trend without us specifying the shape of it. This can be compared to a parametric regression where e.g. a quadratic regression only can find linear and quadratic trends. The small deviation from the real sine-function comes from the largest outliers in the noisy data.

The two lower figures on page 16 demonstrate the importance of having an automatic algorithm for setting the bandwidth. The kernel regression is rather sensitive to the choice of bandwidth and a naive method where the bandwidth is set to a predetermined number might give arbitrary results. If too high a bandwidth is set, too much weight is given to distance observations and the kernel equation gets too linear. This is visualised in the lower leftmost figure where a kernel regression with bandwidth 25 is showed. In the opposite way, too low a bandwidth gives too noisy kernel functions since too little weight is given to distant observations. This is visualised in the lower rightmost figure where a bandwidth of 0.7 is used.

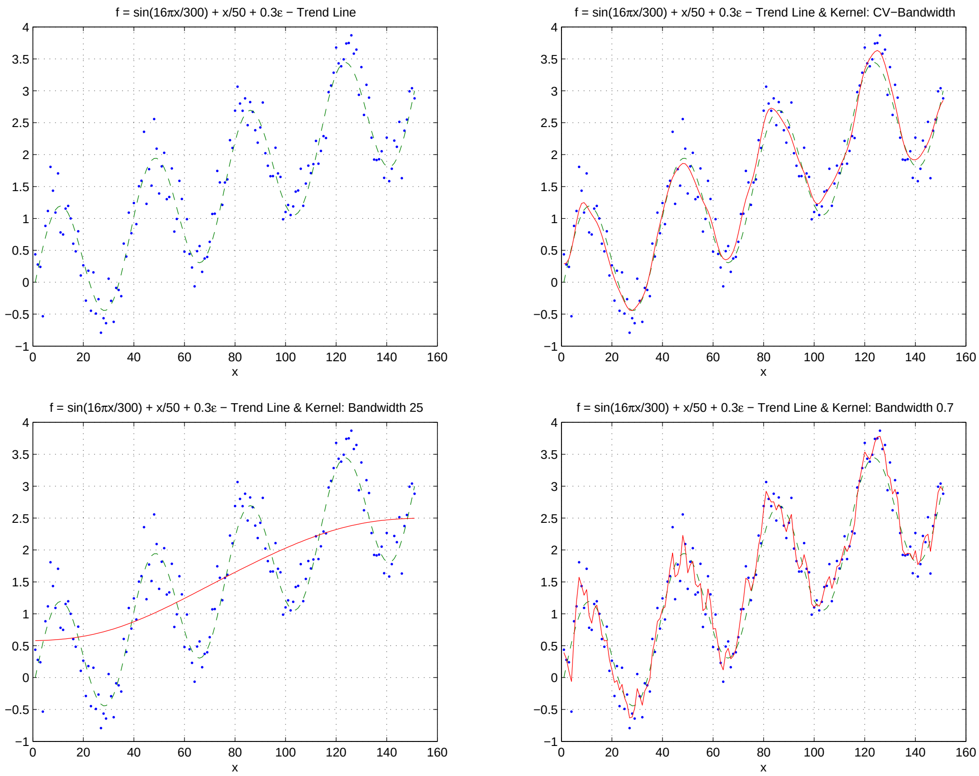

3.4 Statistical Distribution Tests

3.4.1 Goodness-of-Fit Test

A general distribution test, which tests for equality of two distributions independent of distributional assumptions, is the goodness-of-fit test. The test is performed by dividing an unconditional distribution into an arbitrary number of equally-length intervals. The cut-off points between the intervals are used to split the observations in the conditional sample into corresponding intervals. If the conditional and the unconditional samples are drawn from the same distribution the relative frequency of observations in each interval should be equal. More generally, if the unconditional sample is divided into m intervals, the relative frequency of observations in each interval of the conditional sample should be $1/m$. How well the fraction of observations in each interval of the conditional sample fits with fractions of the unconditional sample can be tested both aggregated and on an individual interval level. On an individual interval level, using $m = 10$ intervals, the relative frequency δ_j of unconditional observations in interval j, $j = 1, ..., 10$, are asymptotically distributed:

$$\sqrt{n}(\hat{\delta}_j - 0.10) \sim N(0, 0.10(1 - 0.10))$$

which yields that the t-statistic can be given as:

$$t_j = \frac{\hat{\delta}_j - 0.10}{\sqrt{\frac{0.10(1-0.10)}{n}}} \sim t_n$$

where n is the total number of observations in the unconditional sample. On an aggregated level the corresponding goodness-of-fit test statistic Q is given by:

$$Q = \sum_{j=1}^{10} \frac{(n_j - 0.10n)^2}{0.10n} \sim \chi_9^2$$

where n_j is the number of observations in each of the $j = 1, ..., 10$ intervals and n is the total number of observations in the unconditional sample.

If the relative frequencies of observations in the unconditional sample are close to 0.1, the conditional sample includes no incremental information and the observed t and Q values should be small. If incremental information exists, t and Q should deviate from zero (Lo, Mamaysky and Wang, 2000).

3.4.2 Kolmogorov-Smirnov Two Distributions Test

The Kolmogorov-Smirnov two distributions test is a test of whether two independent samples are drawn from the same distribution. The Kolmogorov-Smirnov test does not require any assumptions about the underlying distribution and is sensible to any kind of difference between the two samples, be it centrality, dispersion of skewness etc. (Castellan Jr and Siegel, 1988). Let $X_1, X_2, ..., X_{n_1}$ and $Y_1, Y_2, ..., Y_{n_2}$ be i.i.d. random variables drawn from a continuously cumulative density function F and G respectively. The Kolmogorov-Smirnov test tests the hypothesis that the two density functions are identical:

$$H_0\colon F(x) = G(x)$$

against the case that the distributions are different:

$$H_1\colon F(x) \neq G(x)$$

If the samples have been drawn from the same distribution the cumulative density functions of the both samples should be close to each other. If the cumulative density functions, at any point, are too far from each other it suggests that the samples are not drawn from the same distributions. For large samples ($n_1, n_2 > 25$) the test statistics is determined by the maximum distance between the two cumulative distribution functions:

$$D_{n_1,n_2} = (\frac{n_1 n_2}{n_1 + n_2})^{1/2} \max_{-\infty<x<\infty} \mid F(x) - G(x) \mid$$

which Smirnov (1939) has shown is asymptotically distributed as:

$$P(D_{n_1,n_2} \leq x) = \sum_{k=-\infty}^{\infty} (-1)^k e^{-2k^2x^2}$$

If the test statistics is larger than the cut-off point the null hypothesis is rejected. It is vital to note that the distribution of the test statistics is discrete in terms of the number of intervals which should be used in the test. If too few intervals are used, information about the cumulative distribution will be lost.

4 Methodology

4.1 Overview of Methodology

On an aggregated level, the Lo, Mamaysky and Wang (2000) methodology, which serves as a base for our study, can be described in a number of distinct steps:

1. A number of actively traded, listed stocks are defined and the price-series for each stock is collected. The full time-series is split into a number of sub-periods to control for potential non-stationarities induced by changing market structure.

2. For each stock and in each sub-period, a series of rolling windows where patterns can be searched for are defined.

3. For each window, a kernel estimator of Gaussian type is constructed and all local extrema are determined. Based on predefined geometrical properties of the local extrema for each pattern, the occurrence of a pattern can be determined.

4. When a pattern is found a one-day continuously compounded return following the pattern is calculated. This yields a sample of returns conditional on the occurrence of patterns. All returns are normalized to enhance comparability between stocks and time-periods.

5. A sample of unconditional one-day continuously compounded normalized returns is created for comparison with the unconditional sample.

6. The information content in the technical patterns is determined by comparing the distribution of the conditional returns with the unconditional distribution of returns using a standard goodness-of-fit test and a Kolmogorov-Smirnov two distributions test. Both tests are testing the null hypothesis that the two samples are drawn from the same distribution. If the patterns include no incremental information the conditional and the unconditional distribution will be similar and the null hypothesis will hold. In contrast, if the patterns actually include incremental information about future stock prices the null hypothesis will be rejected.

4.2 Data

We have chosen to look at data for all stocks now listed on the OMX Swedish Large Caps, giving us 70 different stocks. There are a number of reasons for us using Swedish data. First of all it is our home market and we think it would be more interesting both for us and for our readers to study the Swedish market. We also consider ourselves having general knowledge about the market and the companies, something that can be helpful when conducting the analysis. Further on, the Swedish stock exchange is among the most liquid in the world (see, for example, Eun and Resnick (2007)). Also, companies from many different industries are listed making our results more general. Finally, and most relevant, as far as we are concerned no empirical study of this type has been conducted on Swedish data, making our study contributing to the research community.

We use daily closing prices from the period between January 1, 1982, and December 31, 2006, if available, for the specific stock. If not, we use the first trading date available. The reason for choosing this time-period is simply that it was the longest time-series available with reliable data. Needless to say, a longer time-horizon we believe would have made our results more reliable, but we feel confident in saying that our sample is large enough. As data provider we used the Six Trust database.

We are only analysing stocks currently listed on the OMX Swedish Large Caps list. This means that we are ignoring stocks that have been delisted during the sample period. This could have caused the issue of survivorship bias but this is not a problem in our study. We are only interested in finding the technical patterns for each stock, if it has performed relatively good or bad over the period does not affect our results. There is no reason to believe that delisted stocks have more, less or more informative patterns than stocks which are still listed.

4.3 Preparation of Data

Before starting out the analysis, a number of adjustments to our data set has been made in order to match it to the scope of the analysis. All observations where the traded volume was zero were deleted. By definition, in these cases the price does not fluctuate and the observation is therefore not relevant to our study. Also, a control for liquidity was performed. As liquidity measure, we require the price of a stock to change from one date to the other in more than 75 percent of the observations. If there is low volatility, we believe the patterns will be unreliable and the stock should therefore not be included in our study. This correction leaves us with 65 stocks to analyze. By the same token, for companies that have several classes of stocks listed we picked the most liquid one. Finally, to reduce the effects of nonstationarities resulting from changing market structure and institutions and in line with Lo, Mamaysky and Wang (2000), the sample was divided into five sub-periods of five years: 1982 to 1986, 1987 to 1991, 1992 to 1996, 1997 to 2001 and 2002 to 2006.

4.4 Discrepancies in the Data Sets

We have tried to keep the differences between our data set and that of Lo, Mamaysky and Wang (2000) as irrelevant as possible within the scope of our study. However, some discrepancies have been impossible to avoid. Lo, Mamaysky and Wang (2000) use US stocks from both NYSE and Nasdaq for the time-period between 1962-1996, we use Swedish stock data from OMX Swedish Large Caps for the time-period between 1982-2006. Lo, Mamaysky and Wang (2000) divide their data into 7 sub-periods, 1962-1966, 1967-1971 and so on, we divide our data into 5 sub-periods, 1982-1986, 1987-1991 and so on. Lo, Mamaysky and Wang (2000) sample 50 out of 350 stocks for the 7 time-periods of 5 years, with all observations present at all periods. We use all stocks listed on OMX Swedish Large Caps as of today, and use as long a time-series as we can find in Six Trust. This makes our sample size different for each of the five sub-periods we use. Lo, Mamaysky and Wang (2000) extend their analysis and divide their full sample into five sub-samples based on the size of the firm. We choose not to divide our sample in the similar way, since we believe that the division by time-period is enough to enhance comparability.

Lo, Mamaysky and Wang (2000) use a rule that 75 percent of the observations for a time-series must be non-missing. We use a rule stating that the price has to change between two consecutive days for 75 percent of the observations for the stock to be part of our sample. This deducts 5 stocks from the data set, leaving us with 65 stocks. We believe that these rules are comparable since Six Trust tables an old value instead of a missing one. The differences stated above gives us a final sample of approximately half the size of Lo, Mamaysky and Wang (2000), 207 192 for us versus 423 556 for NYSE and 411 010 for Nasdaq.

Further on, Lo, Mamaysky and Wang (2000) perform one geometric brownian motion and table the results. We feel that this makes the results rather arbitrary, thus we perform five geometric brownian motions. Lo, Mamaysky and Wang (2000) use a sample of 350 stocks for their unconditional sample for the goodness-of-fit test. We do not have that many stocks so we use the full sample as the unconditional sample for the goodness-of-fit test.

4.5 Defining Windows

Let each stock in our data set represent a series of prices $\{P_1, ..., P_T\}$ where T is the number length of the time-series corresponding to the given stock. We divide each series of prices into windows of length $l+d$ on a rolling basis from time $t = 1$ to $t = T - l - d + 1$. The parameter l represents the number of historical data points used to detect a pattern while d is the number of days between the completion of the pattern and the first day on which the pattern can be acted on. Lo, Mamaysky and Wang (2000) set $l = 35$ to limit themselves to short-term patterns and they assume it takes $d = 3$ days to recognize a pattern after it has been completed. We used the same numbers to enhance comparability, thus each window consists of 38 observations.

4.6 Kernel Regression

For each sub-window we estimate a kernel regression using the prices in that window as:

$$\hat{m}_k(\tau) = \frac{\sum_{s=t}^{t+l+d-1} K_h(\tau - s)P_s}{\sum_{s=t}^{t+l+d-1} K_h(\tau - s)} \ , \ t = 1, ..., T - l - d + 1$$

where τ represents each observation in the sub-window. The bandwidth is set to the bandwidth which minimizes the cross-validation function $CV(h)$ with one observation at the time omitted:

$$CV(h) = \frac{1}{T} \sum_{s=t}^{t+l+d-1} (P_s - \hat{m}_{h,t})^2 \ , \ t = 1, ..., T - l - d + 1$$

Lo, Mamaysky and Wang (2000) confront a number of professional technical analysts with the kernel estimations produced with the optimal bandwidth from the cross-validation function and conclude that the bandwidth is slightly too high to capture the full information used by the technical analysts. They therefore multiply the optimal bandwidth h^* with 0.3. The derivation of the bandwidth multiplier is done before the analysis and the derived bandwidth is kept throughout the whole study to prevent data mining (Lo, Mamaysky and Wang, 2000). This choice can be questioned as being slightly ad hoc from an analytical point of view. However, it is out of the scope for this thesis to question the tacit knowledge of these technical analysts and we therefore use the same bandwidth multiplier in our study. Furthermore, this enhances one of the purposes of our study; to conduct an out of sample test of the Lo, Mamaysky and Wang (2000) methodology.

The procedure yields an estimate of $m_h(\tau)$ which is a differentiable function of τ. Local extrema are then identified by finding τ such that $Sgn(m_h'(\tau)) =$

$-Sgn(m'_h(\tau+1))$ (Lo, Mamaysky and Wang, 2000). Once such τ is detected we look for the corresponding extrema in the interval $[\tau-1, \tau+1]$ in the original price-series $\{P_t\}$. Following the Lo, Mamaysky and Wang (2000) methodology we use the next observation where $m'_h(\tau) \neq 0$ as a base of comparison if we detect a τ where $m'_h(\tau) = 0$, which occurs if the closing price stays the same for several days.

4.7 Defining Patterns

Once we have detected all local extrema in a window we look in the price-series $\{P_t\}$ for the patterns predefined by Lo, Mamaysky and Wang (2000).

Definition 1: *Head-and-shoulders (HS) and inverted head-and-shoulders (IHS) patterns are characterized by a sequence of five consecutive local extrema, $E_1, ..., E_5$, located such that:*

$$HS = \begin{cases} E_1 \text{ a maximum} \\ E_3 > E_1,\ E_3 > E_5 \\ E_1 \text{ and } E_5 \text{ within 1.5 percent of their average} \\ E_2 \text{ and } E_4 \text{ within 1.5 percent of their average} \end{cases}$$

$$IHS = \begin{cases} E_1 \text{ a minimum} \\ E_3 < E_1,\ E_3 < E_5 \\ E_1 \text{ and } E_5 \text{ within 1.5 percent of their average} \\ E_2 \text{ and } E_4 \text{ within 1.5 percent of their average} \end{cases}$$

Definition 2: *Broadening tops (BTOP) and bottoms (BBOT) are characterized by a sequence of five consecutive local extrema $E_1, ..., E_5$ such that:*

$$BTOP = \begin{cases} E_1 \text{ a maximum} \\ E_1 < E_3 < E_5 \\ E_2 > E_4 \end{cases}$$

$$BBOT = \begin{cases} E_1 \text{ a minimum} \\ E_1 > E_3 > E_5 \\ E_2 < E_4 \end{cases}$$

Definition 3: *Triangle tops (TTOP) and bottoms (TBOT) are characterized by a sequence of five consecutive local extrema $E_1, ..., E_5$ such that:*

$$TTOP = \begin{cases} E_1 \text{ a maximum} \\ E_1 > E_3 > E_5 \\ E_2 < E_4 \end{cases}$$

$$TBOT = \begin{cases} E_1 \text{ a minimum} \\ E_1 < E_3 < E_5 \\ E_2 > E_4 \end{cases}$$

Definition 4: *Rectangle tops (RTOP) and bottoms (RBOT) are characterized by a sequence of five consecutive local extrema* $E_1, ..., E_5$ *such that:*

$$RTOP = \begin{cases} E_1 \text{ a maximum} \\ \text{Tops within 0.75 percent of their average} \\ \text{Bottoms within 0.75 percent of their average} \\ \text{Lowest top } > \text{ Highest bottom} \end{cases}$$

$$RBOT = \begin{cases} E_1 \text{ a minimum} \\ \text{Tops within 0.75 percent of their average} \\ \text{Bottoms within 0.75 percent of their average} \\ \text{Lowest top } > \text{ Highest bottom} \end{cases}$$

Definition 5: *Double tops (DTOP) and bottoms (DBOT) are characterized by an initial local extremum* E_1 *and a subsequent local extrema* E_a *and* E_b *such that:*

$$E_a = \max\{P_{t_k^*} \ : \ t_k^* > t_1^* \, , k = 2, ..., n\}$$
$$E_b = \min\{P_{t_k^*} \ : \ t_k^* > t_1^* \, , k = 2, ..., n\}$$

and

$$DTOP = \begin{cases} E_1 \text{ a maximum} \\ E_1 \text{ and } E_a \text{ within 1.5 percent of their average} \\ t_a^* - t_1^* > 22 \end{cases}$$

$$DBOT = \begin{cases} E_1 \text{ a minimum} \\ E_1 \text{ and } E_b \text{ within 1.5 percent of their average} \\ t_b^* - t_1^* > 22 \end{cases}$$

where t_1, t_a and t_b are the times for the local extremas E_1, E_a and E_b.

4.8 Conditioning on Volume

As stated in the theory part, the trading volume plays an important role in technical analysis. Lo, Mamaysky and Wang (2000) use a rather primitive method when controlling for changing volume. To make our results comparable we use the same method. For each stock in each window they calculate the average trading volume during the first and the second halves of each window. If we denote the average volume in the first period υ_1 and in the second period υ_2 the following conditions are made. If $\upsilon_1 > 1.2\upsilon_2$ it is categorized as a decreasing volume event. In the same way, if $\upsilon_2 > 1.2\upsilon_1$ it is categorized as an increasing volume event. If neither inequality holds, the volume is said to be neutral.

4.9 Calculating Returns

If we find a pattern completed at time $t + l + 1$, using information from time t to $t + l + d - 1$, we compute a one day continuously compounded return $R = ln(P_t + l + d + 1/P_t + l + d)$. Note that no forward looking information is used and the return is completely out of sample and do not suffer from look-ahead bias (Lo, Mamaysky and Wang, 2000). Since the different stocks differs quite a lot in terms of expected return and standard deviation we normalize the returns to make them more comparable between stocks:

$$X_{it} = \frac{R_{it} - \mathbf{E}[R_{it}]}{\sigma[R_{it}]}$$

where $\mathbf{E}[R_{it}]$ and $\sigma[R_{it}]$ are the mean and the standard deviation of the return of each stock in each five year sub-period respectively. As a consequence of the normalization, each return has zero mean and unit standard deviation (Lo, Mamaysky and Wang, 2000). This is a necessary condition when using a goodness-of-fit test to compare the distributions.

4.10 Creating a Sample for Comparison

To compare with the conditional returns, we create a sample of unconditional returns following the procedure of Lo, Mamaysky and Wang (2000). For each stock and for each of the five sub-periods, as well as for the whole time-period, we compute one-day continuously compounded returns using non-overlapping intervals. As for the conditional returns, the unconditional returns are normalized by deducting the mean and dividing by the standard deviation calculated for each stock in each sub-period. The returns from all stocks are then combined into a single sample for each sub-period. This procedure yields returns with zero mean and unit variance, which distributions can be compared with the conditional returns.

4.11 Comparing the Distributions of the Returns

The conditional returns are compared to the unconditional returns using a goodness-of-fit test and a Kolmogorov-Smirnov two distributions test. Ideally, one would like to determine if these technical patterns provide information which could be used to find profitable trades. However, to determine if the technical patterns more than compensate for the risk involved in exploring them one has to compare the conditional returns with the fair returns given both the syncratic and the idiosyncratic risk. Lo, Mamaysky and Wang (2000) conclude that, in absence of a fully specified dynamic general equilibrium asset-pricing model, such a comparison can not be made. We believe that no convincing evidence that such model exists is currently available. Trying to compare the conditional returns to any currently existing pricing model would only add yet another layer of uncertainty to the validity of the results. We therefore follow the Lo, Mamaysky and Wang (2000) method and limit ourselves to examine the information contents in these patterns. If no incremental information exists in these patterns, the conditional and the unconditional returns should be drawn from the same distribution. Oppositely, if the distributions differ, the patterns can be said to include incremental information. The similarity of the conditional and the unconditional return distributions can be tested using a goodness-of-fit test and a Kolmogorov-Smirnov two distributions test. If the null hypotheses, as described in the theory section, can be rejected, incremental information exists in these patterns.

4.12 Graphical Representation

The figure on page 27 showes a graphical illustration of our algorithm detecting a broadening tops pattern for the Ericsson stock[1].

[1]For a graphical representation of each individual pattern determined by our detection algorithm, see Appendix II, pages 43-47.

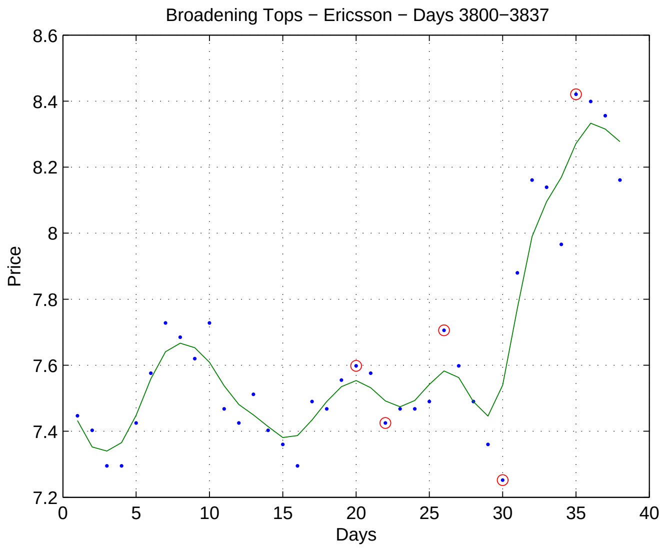

The figure shows a kernel regression constructed with data from a 38-day window. Nine local extrema can be found in the underlying trend. For each extrema in the kernel estimate the corresponding extrema of the price function is determined in the interval [-1,+1] days. Five of these price-series extrema, marked with circles, form a broadening tops pattern. As the final maxima in the price-series occurs at day 35 in the window, a one day continuously compounded return is calculated between the last observation in the window and the first after the window. Thus, as described, no forward-looking information is used.

4.13 Comparing with Simulated Geometric Brownian Motions

It is crucial to understand whether the detected patterns seem to be causal or if they are nothing but effects of a totally random process. For each stock, we create a fictive price-series where the price function follows a stochastic process. In particular, we let the price functions follow a geometric brownian motion:

$$dP_t = \mu P_t dt + \sigma P_t dW_t$$

where W_t is a standard brownian motion with the drift μ and the diffusion σ estimated from the real price-series for each individual stock. To make the data completely comparable we use the same length and starting value P_0 for each simulated geometric brownian motion as for the corresponding stock.

When simulating stochastic processes each simulation is unique. A comparison with a single geometric brownian motion thus risks being fairly arbitrary. We therefore simulate 5 individual processes for each stock and compare the real price-series with the simulations. If the true patterns are occurring for no particular reason we expect the results from the geometric brownian motion to mirror the results from the real price-series.

5 Analysis and Results

5.1 Frequencies

In the table on page 37 the total number of observations, the number of the different patterns discovered and the fraction of the observations included in a pattern, across the whole sample and across each sub-period, are displayed. The most common patterns in our sample are the head-and-shoulders and the inverted head-and-shoulders, with 1 115 and 1 076 occurrences respectively. The second most frequent patterns are the double tops and bottoms. The rest of the patterns are distributed quite equally among themselves.

Across the sub-periods, the relative frequency of each pattern in each sub-period corresponds rather closely to the total number of patterns for that sub-period. The biggest discrepancies occur in the last two sub-periods where the patterns occur more frequent. A hypothesis could be that these two periods have been the most volatile giving a larger chance for these patterns to occur. We also find slightly more patterns when conditioning on decreasing volume than when conditioning on increasing volume. The distribution among the patterns is fairly similar when conditioning on volume, though.

For the purpose of our study, we also compare our results to the findings of Lo, Mamaysky and Wang (2000). In order to do that, we need to take into consideration that we have approximately half the number of observations in our sample, and that our observations are not equally distributed across the sub-periods. The table on page 37 shows that the fraction of observations where we can find a pattern is about the same, 3.0 percent in our results and 3.4 percent for NYSE and 2.2 percent for Nasdaq in Lo, Mamaysky and Wang (2000). This could be seen as a rough estimate that the occurrences of patterns are as likely in both American as well as in Swedish stock data. Also, we see that the distribution among the different patterns is slightly reversed compared to our data. For NYSE, the most common patterns are the double tops and bottoms followed by the head-and-shoulders and the rectangles. For Nasdaq, the rectangles are the most common followed by the doubles and the head-and-shoulders. One possible hypothesis could be that our sample contains the volatile period between 1997 and 2006 which is absent in the sample of Lo, Mamaysky and Wang (2000). The patterns in Lo, Mamaysky and Wang (2000) are also relatively equally distributed over the sub-periods, which implies a slight difference from our results. As a conclusion, however, we would still argue that our results do not differ significantly from Lo, Mamaysky and Wang (2000).

5.2 Geometric Brownian Motions

For purpose of comparison, we have also simulated five individual runs of geometric brownian motions for each stock in our study. For each geometric brownian motion sample we have run the algorithm for pattern recognition. The table on page 38 shows that the random walk models all generate similar results in terms of number of patterns. The fractions of patterns are 2.4 percent or 2.5 percent for the geometric brownian motion samples compared to 3.0 percent for our original sample. Furthermore, we see that the distributions of the patterns are more even in the geometric brownian motion samples, compared to high peaks for the head-and-shoulders, the inverted head-and-shoulders and the doubles, for our original sample. These results from the simulations indicate that there is a difference between our stock data and independently and identically distributed lognormal returns. This suggests that the occurrences of patterns are "there for a reason" and not an effect of a random process.

5.3 Descriptive Statistics

The table on pages 39-40 reports the descriptive statistics - mean, standard deviation, skewness and kurtosis - of unconditional and conditional normalized one-day returns for our sample. The column named raw provides data for all the observations across our sample and the other columns provides data for each pattern respectively. Our statistics, as well as Lo, Mamaysky and Wang (2000), show quite significant differences in the normalized return distributions. The first four moments of normalized raw returns are 0.000, 1.000, -0.1481 and 18.2433. The same statistics for e.g. inverted head-and-shoulders are 0.039, 0.8266, -0.4686 and 7.0239. Most other patterns show any other discrepancies[2]. All these statistics for the ten conditional patterns, and the variances in the conditional and unconditional return populations, suggest that identifying and using the ten technical patterns does have some effect on the distribution of returns. Thus, this is a first sign that incremental information might be present in these conditional samples.

[2]These numbers should, however, be interpreted with a bit of caution since some of them are largely affected by outliers. The extremely high kurtosis for TTOP and DBOT are explained by the drop of Alfa Laval from SEK 66.5 to SEK 45.3 on November 1, 2002, and the drop of Axfood from SEK 220 to SEK 181.5 on January 27, 2005. Without these observations, the kurtosis is more normal, 4.5560 and 7.6864 respectively. Fortunately, the goodness-of-fit and the Kolmogorov-Smirnov tests are much less affected by single outliers.

5.4 Goodness-of-Fit Test

The table on page 41 reports the diagnostics of the goodness-of-fit test for our sample. The table is, conditioned on each pattern, divided into deciles with each decile containing 10 percent of the observations of the unconditional return population and the last column showing the goodness-of-fit test statistic Q and corresponding p-values. If conditioning on the pattern provides no information, the expected percentage falling into each decile is 10 percent. The table shows that our sample has very low t-values overall. Furthermore, only three out of ten patterns - TTOP, IHS and RBOT - have relative frequencies of the conditional returns that are significantly different from those of the unconditional returns on the 5 percent level. On the 10 percent level, three more patterns are added. The conclusion is that these goodness-of-fit results show weak support for the informative content of the technical patterns compared to the results of Lo, Mamaysky and Wang (2000). In Lo, Mamaysky and Wang (2000), the NYSE sample gives seven patterns showing statistically significant differences whereas for Nasdaq all patterns are significantly different. One plausible explanation in defense of technical analysis is that our sample includes too few observations to generate reliable results for a goodness-of-fit test. For example, when comparing our results for BBOT with those of Lo, Mamaysky and Wang (2000), one can see that the percentage of observations in each decile deviates quite a lot from the expected 10 percent in our test. Unfortunately, the limited number of conditional returns compared to Lo, Mamaysky and Wang (2000) makes the significanses weak.

5.5 Kolmogorov-Smirnov Two Distributions Test

In the table on page 42 we tabulate the p-values of the results of the Kolmogorov-Smirnov two distributions test of the equality of the conditional and unconditional distributions. The results are shown for the full sample size and the five sub-periods, divided into all observations, conditioned on increasing volume, conditioned on decreasing volume and the difference between the increasing and decreasing volume-trend distributions.

When looking at the result for all observations over the full sample period, five of the ten patterns - BTOP, TTOP, HS, IHS and RBOT - are statistically significant on the 5 percent level. On the 6 percent level, RTOP is also added. These patterns show p-values from 0.0001 for TTOP to 0.0562 for RTOP. On the other hand, the other patterns show p-values ranging from 0.2182 to 0.3689. These results indicate that there is some incremental information in many of the technical indicators.

When looking into the sub-periods, the results are more varying. The first three periods, ranging from 1982 to 1996, are difficult to interpret because of the small sample sizes. The Kolmogorov-Smirnov test requires quite a large number of observations to be reliable. For the period between 1997 and 2001 five of the ten patterns are statistically significant, whereas between 2002 and 2006, only one is so. But yet again, the results might be arbitrary due to a limited sample size.

5.6 Conditioning on Volume

When conditioning on volume, both increasing and decreasing, the table on page 42 showes that the statistical significance increases overall for most of the patterns. When looking at the full sample, seven out of ten patterns are statistically significant for both increasing and decreasing volume. However, it is not the same patterns in both cases. Surprisingly, when examining the sub-periods, we see that the last period, between 2002 and 2006, now show three and four statistically significant patterns for the increasing and decreasing condition respectively. This is to be compared with only one for the full sample. This indicates that there is some information in the volume. The difference between the increasing and decreasing volume-trend conditional distributions is statistically insignificant for all the patterns, though. This can partly be explained by the fact that the sample size is fairly small reducing the power of the Kolmogorov-Smirnov test, but it is also an indication of the fact that we cannot differentiate between increasing and decreasing volume. Thus, volume seems to matter, but perhaps not necessarily in the way defined by Lo, Mamaysky and Wang (2000).

6 Concluding Remarks

We have performed an out-of-sample test of the Lo, Mamaysky and Wang (2000) methodology for analysis of technical indicators and at the same time evaluated the informative content of these indicators on Swedish stock data.

On the whole, our results are fairly similar to those of Lo, Mamaysky and Wang (2000), even though a completely different data set is used. We find just about as many patterns in comparison to the number of observations as they do, and the distributions of the patterns are comparable. Neither when conditioning on different time-periods nor on volume trends the results deviate substantially from those of Lo, Mamaysky and Wang (2000). Furthermore, the geometric brownian motion simulations indicate that the patterns found are at least not stochastic. These results indicate that the Lo, Mamaysky and Wang (2000) model does a good job in finding the technical indicators also on Swedish data. The model appears to be fairly general and not over-fitted to their data set. Our study is a sign of the fact that the algorithm might serve as an analytical tool also for other sets of data.

When assessing the question of informative content in technical indicators in Swedish stock data, our study shows that these patterns might indeed contain some incremental information. There are differences in all four moments between the conditional and unconditional samples of normalized returns for all time-periods, which might suggest that new information can be extracted from those patterns. In addition, the signs of the means are analogue with the conventional theory on the subject. Furthermore, the Kolmogorov-Smirnov test yields significant results for six out of ten patterns at the 6 percent level. The results get even more significant when conditioning on volume, also in line with theory which emphasize the importance of trading activity. The only discrepancy between our results and those of Lo, Mamaysky and Wang (2000) is the goodness-of-fit test where Lo, Mamaysky and Wang (2000) get slightly more significant results. However, the difference in statistical significance is most likely due to our smaller sample size.

It should be noted that our results only show that the technical indicators provides incremental information, not that technical analysis by default can generate excess trading profits. To test for this we need a general equilibrium pricing model, which not only prices the idiosyncratic risk, but rather the full risk an investor is exposed to when holding an undifferentiated portfolio. As no such model is currently available we limit ourselves to conclude that these patterns contain additional information which could be included in an investment strategy. In any case, this incremental information is a contradiction of the weak form of the efficient market hypothesis, which states that future returns can not be explained by past returns.

References

[1] Baltagi, Badi H., 2001, A Companion to Theoretical Econometrics, 1st Edition. Oxford, UK: Blackwell Publishers Ltd.

[2] Basu, S, 1977, Investment Performance of Common Stocks in Relation to their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis, *The Journal of Finance* 32: 663-682.

[3] Blume, Lawrence, Easley, David and Maureen O'Hara, 1994, Market Statistics and Technical Analysis: The Role of Volume, *Journal of Finance* 49: 153-181.

[4] Brock, William, Lakonishok, Joseph and Blake LeBaron, 1992, Simple Technical Trading Rules and the Stochastic Properties of Stock Returns, *Journal of Finance* 47: 1731-1764.

[5] Brown, David and Robert Jennings, 1989, On Technical Analysis, *Review of Financial Studies* 2: 527-551.

[6] Campbell, John Y., Lo, Andrew W. and A. Craig MacKinlay, 1997, The Econometrics of Financial Markets, 2nd Edition. Princeton, NJ: Princeton University Press.

[7] Castellan Jr, N. John and Sidney Siegel, 1988, Nonparametric Statistics for the Behavioural Sciences, 2nd Edition. Singapore: McGraw-Hill.

[8] Chan, Louis K.C., Jegadeesh, Narasimhan and Josef Lakonishok, 1996, Momentum Strategies, *The Journal of Finance* 51: 1681-1713.

[9] DeBondt, Werner F.M. and Richard Thaler, 1985, Does the Stock Market Overreact?, *Journal of Finance* 40: 793-805.

[10] DeBondt, Werner F.M. and Richard Thaler, 1987, Further Evidence of Investor Overreaction and Stock Market Seasonality, *Journal of Finance*, 42: 557-581.

[11] Dittmar, Robert, Neely, Christopher and Peter Weller, 1997, Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach, *Journal of Financial and Quantitative Analysis* 32: 405-426.

[12] Edwards, Robert and John Magee, 1997, Technical Analysis of Stock Trends, 7th Edition. Chicago, Ill: John Magee Inc.

[13] Eun, Cheol S. and Bruce G. Resnick, 2007, International Financial Management, 4th Edition. New York, NY: McGraw-Hill/Irwin.

[14] Fama, E.F. and K.R. French, 1992, The Cross-Section of Expected Stock Returns, *The Journal of Finance* 47: 427-465.

[15] Fama, E.F., 1991, Efficient Capital Markets: II, *The Journal of Finance* 46: 575-1617.

[16] Fan, J. and I. Gijbels, 1996, Local Polynomial Modelling and Its Applications, 1st Edition. Suffolk, UK: Chapman & Hall.

[17] Green, P.J. and B.W. Silverman, 1994, Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach, 1st Edition. London, UK: Chapman & Hall.

[18] Härdle, Wolfgang, 1990, Applied Nonparametric Regression, 1st Edition. London, UK: Press Syndicate of the University of Cambridge.

[19] Jegadeesh, N. and S. Titman, 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *The Journal of Finance* 48: 65-91.

[20] Lakonishok, J. and S. Smidt, 1988, Are Seasonal Anomalies Real? A Ninety-Year Perspective, *Review of Financial Studies* 1: 403-425.

[21] Lo, Andrew W, Mamaysky, Harry and Jiang Wang, 2000, Foundations on Technical Analysis: Computational Algorithms, Statistical Inference, and Empirical Implementation, *The Journal of Finance* 55: 1705-1765.

[22] Lo, Andrew W. and A. Craig MacKinlay, 1988, Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies* 1: 41-66.

[23] Lo, Andrew W. and A.Craig MacKinlay, 1999, A Non-Random Walk Down Wall Street, 5th Edition. Princeton, NJ: Princeton University Press.

[24] Mittelhammer, Ron C., Judge, George G. and Douglas J. Miller, 2000, Econometric Foundations, 1st Edition. London, UK: Cambridge University Press.

[25] Neftci, Salih, 1991, Naive Trading Rules in Financial Markets and Wiener-Kolmogorov Prediction Theory: A Study of Technical Analysis, *Journal of Business* 64: 549-571.

[26] Nilsson, Peter and Johnny Torssell, 2000, Boken om teknisk analys - Teori, grunder och tillämpning, 1st Edition. Stockholm, Swe: Börsinsikt AB.

[27] Pring, M.J., 1980, Technical Analysis Explained, 1st Edition. New York, NY: McGraw-Hill.

[28] Pruitt, Stephen and Robert White, 1988, The CRISMA Trading System: Who Says Technical Analysis can't Beat the Market?, *Journal of Portfolio Management* 14: 55-58.

[29] Rhea, R, 1932, Dow Theory. New York, NY: Barron's.

[30] Rouwenhorst, Geert, 1998, International Momentum Strategies, *Journal of Finance* 53: 267-284.

[31] Schwert, G., 2003, Anomalies and Market Efficiency, Chapter 17 of the Handbook of Economics and Finance (G. Constantinides, M. Harris and R. Stulz, eds.). Amsterdam, Ne: North Holland.

[32] Smirnov, N., 1939, On the Estimation of the Discrepancy between Empirical Curves of Distribution for two Independent Samples, *Bulletin. Math. Univ. Moscow* 2, 3-14.

[33] Treynor, Jack and Robert Ferguson, 1984, In Defence of Technical Analysis, *Journal of Finance* 40: 757-773.

A Appendix I - Tables

A.1 Frequencies

The table shows the total number of recognized patterns for each of our ten defined technical indicators found in a sample of 65 stocks from OMX Swedish Large Caps from 1982 to 2006. As a benchmark, the same kind of dataset from Lo, Mamaysky and Wang (2000) is given (LMW NYSE and LMW Nasdaq). The sample is divided into the full sample and 5 sub-periods. The column named "Sample" shows whether the frequency counts are unconditional on volume (Full Sample, 1982-1986 etc.), conditioned on decreasing volume (Dec. Vol.) or conditioned on increasing volume (Inc. Vol.). The fraction column (Frac.) indicates the percentage of all observations included in a pattern across the whole sample and across each of the five sub-periods.

Sample	Raw	HS	IHS	BTOP	BBOT	TTOP	TBOT	RTOP	RBOT	DTOP	DBOT	Frac.
LMW NYSE	423 556	1294	1193	1611	1654	725	748	1482	1616	2076	2075	3.4%
LMW Nasdaq	411 010	919	817	414	508	850	789	1134	1320	1208	1147	2.2%
Full Sample	207 192	1115	1076	412	337	478	533	437	516	666	678	3.0%
Dec. Vol.		499	407	115	111	212	219	209	244	226	304	
Inc. Vol.		259	295	153	132	122	142	113	113	218	172	
1982-1986	16 082	68	78	19	20	29	35	27	52	53	59	2.7%
Dec. Vol.		33	42	7	11	12	14	17	38	23	29	
Inc. Vol.		18	17	9	6	9	11	8	8	17	16	
1987-1991	22 946	99	114	40	33	45	56	40	57	70	53	2.6%
Dec. Vol.		48	46	14	12	19	30	22	26	26	28	
Inc. Vol.		22	40	12	14	13	14	11	14	29	16	
1992-1996	36 622	180	159	74	46	82	93	60	61	116	102	2.7%
Dec. Vol.		87	73	26	10	41	42	32	31	43	43	
Inc. Vol.		37	42	33	19	21	27	15	14	34	27	
1997-2001	57 696	337	318	119	122	176	161	96	137	202	211	3.3%
Dec. Vol.		157	102	24	41	81	59	42	64	57	109	
Inc. Vol.		70	94	35	51	40	37	26	34	76	40	
2002-2006	73 846	431	407	160	116	146	188	214	209	225	253	3.2%
Dec. Vol.		174	144	44	37	59	74	96	85	77	95	
Inc. Vol.		112	102	64	42	39	53	53	43	62	73	

A.2 Geometric Brownian Motions

The table shows five simulated geometric brownian motions (GBM1-GBM5), all based on the sample of our study. For each stock, the starting value, the mean and the standard deviation has been the input variables. Shown is also a benchmark from our full sample from section A.1 (Full Sample) and from Lo, Mamaysky and Wang (2000) (LMW NYSE and LMW Nasdaq). Tabulated are the total numbers of recognized patterns for each technical indicator. The fraction column (Frac.) indicates the number of all observations included in a pattern.

Sample	Raw	HS	IHS	BTOP	BBOT	TTOP	TBOT	RTOP	RBOT	DTOP	DBOT	Frac.
LMW NYSE	423 556	1294	1193	1611	1654	725	748	1482	1616	2076	2076	3.4%
GBM NYSE	423 556	1049	1176	577	578	1227	1028	122	113	552	535	1.6%
LMW Nasdaq	411 010	919	817	414	508	850	789	1134	1320	1208	1147	2.2%
GBM Nasdaq	411 010	434	447	1297	1139	1169	1309	96	91	567	579	1.7%
Full Sample	207 192	1115	1076	412	337	478	533	437	516	666	678	3.0%
GBM1	207 192	822	886	463	400	398	449	183	231	560	527	2.4%
GBM2	207 192	790	804	512	419	403	391	211	210	572	600	2.4%
GBM3	207 192	849	814	494	400	423	459	202	244	559	595	2.4%
GBM4	207 192	847	829	480	417	441	484	182	205	623	573	2.5%
GBM5	207 192	860	810	472	377	415	428	218	204	584	572	2.4%

A.3 Descriptive Statistics

The table shows descriptive statistics (mean, standard deviation, skewness and kurtosis) for the 1-day normalised returns of the full unconditional sample (Raw) and conditioned on the ten defined patterns. The sample consists of 65 stocks from OMX Swedish Large Caps from 1982 to 2006. The sample is divided into the full sample and 5 sub-periods. In line with Lo, Mamaysky and Wang (2000), conditional returns of the ten defined patterns are defined as the daily return three days following the conclusion of an occurrence of one of the ten technical patterns. The normalization process of the returns is conducted by subtracting the means and divide by the standard deviations.

Sample	Raw	HS	IHS	BTOP	BBOT	TTOP	TBOT	RTOP	RBOT	DTOP	DBOT
						Full Sample					
Mean	0.0000	-0.0709	0.0390	-0.0608	-0.0617	-0.1682	0.0048	-0.0430	0.0744	-0.0356	-0.0414
Std. Dev	1.0000	0.8831	0.8266	0.8692	1.1055	1.1719	0.9007	0.8391	0.7274	0.9252	1.0388
Skewness	-0.1481	-0.2628	-0.4686	0.0133	-0.5789	-5.4760	-0.1288	0.0219	0.1675	-0.7241	-2.5188
Kurtosis	18.2433	6.9765	7.0239	4.5543	7.9793	77.5185	4.8845	7.1232	6.3053	7.8509	31.6402
						1982-1986					
Mean	0.0000	-0.0511	0.1009	-0.1447	-0.2835	0.0151	0.0219	0.0486	0.3016	0.0294	-0.2137
Std. Dev	1.0000	0.8523	0.8774	1.2608	0.9787	0.8539	1.1057	0.6478	0.7111	1.0321	1.0839
Skewness	-0.0801	-0.4083	0.0059	1.5095	-0.9269	-0.6090	-0.5421	-0.8383	0.2355	-1.6248	-1.0319
Kurtosis	6.5063	4.8393	3.3729	5.9189	3.6789	5.3972	3.9733	5.1825	3.0233	10.3204	5.0960
						1987-1992					
Mean	-0.0000	-0.0937	0.1319	0.1231	0.0598	-0.1814	0.0978	0.1077	-0.1660	-0.1855	-0.0171
Std. Dev	1.0000	0.8559	0.7756	0.7603	1.1504	1.2739	0.7933	0.8220	0.7350	1.0825	0.6646
Skewness	-0.3484	-0.6218	-0.5333	0.0212	0.0468	-0.6232	-0.6436	0.5665	0.2755	-0.4522	-0.2720
Kurtosis	10.5984	4.5684	6.1924	3.7926	3.4262	6.8661	3.9835	3.2622	4.2360	6.4858	3.5456
						1992-1996					
Mean	-0.0000	-0.0928	0.0564	0.1838	-0.4380	-0.1179	-0.1157	-0.0014	0.0822	0.0075	0.0109
Std. Dev	1.0000	0.7801	0.7490	0.9481	1.5235	0.7823	0.9730	0.9072	0.7309	0.9583	0.8773
Skewness	-0.0646	-0.1728	-0.4935	-0.3353	-1.6385	0.7991	-0.3291	1.6773	-0.1223	-1.4748	-0.6390
Kurtosis	15.6000	9.9730	7.4702	4.4081	6.7315	7.5539	4.7650	7.2986	2.8255	14.5331	8.2792

Descriptive Statistics, Continued

Sample	Raw	HS	IHS	BTOP	BBOT	TTOP	TBOT	RTOP	RBOT	DTOP	DBOT
						1997-2001					
Mean	0.0000	-0.1358	0.0153	-0.2330	-0.0387	-0.2055	0.0845	-0.1608	0.0604	0.0950	-0.1140
Std. Dev	1.0000	0.8872	0.9164	0.8333	0.9168	0.9264	0.8765	0.8128	0.7169	0.8801	1.0720
Skewness	0.1166	-0.0374	-0.4081	-0.4115	0.4899	-0.0850	0.2080	-1.1512	0.4966	-0.0680	-3.1671
Kurtosis	9.4805	6.2305	7.5454	4.0511	4.4678	4.7735	3.4881	8.2002	3.9204	4.4594	36.9127
						2002-2006					
Mean	0.0000	-0.0089	0.0129	-0.0818	0.0671	-0.1838	-0.0347	-0.0416	0.0903	-0.1438	0.0333
Std. Dev	1.0000	0.9298	0.7858	0.8021	1.0794	1.5801	0.8722	0.8554	0.7220	0.8547	1.1194
Skewness	-0.4334	-0.3817	-0.6136	-0.0780	0.1531	-7.2523	0.0379	-0.1128	0.0294	-0.6295	-2.7384
Kurtosis	15.4911	7.4148	6.8158	3.8370	8.0048	75.3422	6.3424	6.8618	10.6193	5.3476	34.9282

A.4 Goodness-of-Fit Test

The table shows the results of the goodness-of-fit test for the 1-day conditional normalized returns of our ten defined technical patterns. The sample consists of 65 stocks from OMX Swedish Large Caps from 1982 to 2006. For each pattern, the percentage of conditional returns that fall within each of the ten unconditional-return deciles is tabulated. We should expect the percentage falling into each decile being 10 percent if the pattern condition provides no added value. Asymptotic t-statistics for this null hypothesis are reported in parentheses, and the goodness-of-fit test statistic Q is reported in the last column with the p-value in parenthesis below the statistic. Two asterisks indicates significance on the 5 percent level whereas one asterisk indicates significance on the 10 percent level.

Quantile	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**Q/(p)**
HS	10.22	10.49	11.21	10.04	10.40	11.39	10.13	9.96	7.80	8.34	12.78
	(0.247)	(0.538)	(1.281)	(0.050)	(0.441)	(1.461)	(0.149)	(-0.050)	(-2.736)	(-2.004)	(0.173)
IHS	6.78	8.74	10.22	11.80	11.52	8.74	11.06	10.87	11.15	9.11	24.93
	(-4.194)	(-1.468)	(0.242)	(1.833)	(1.566)	(-1.468)	(1.108)	(0.921)	(1.201)	(-1.017)	(0.003)**
BTOP	11.89	10.68	11.41	11.17	7.52	7.77	10.68	9.47	10.92	8.50	9.21
	(1.187)	(0.447)	(0.899)	(0.751)	(-1.905)	(-1.693)	(0.447)	(-0.370)	(0.600)	(-1.096)	(0.420)
BBOT	13.35	10.09	6.82	8.61	8.01	12.76	12.46	8.01	10.68	9.20	15.49
	(1.810)	(0.054)	(-2.311)	(-0.913)	(-1.344)	(1.518)	(1.369)	(-1.344)	(0.406)	(-0.509)	(0.078)*
TTOP	13.39	10.88	10.88	10.04	11.92	10.25	6.28	8.79	10.88	6.69	20.95
	(2.176)	(0.617)	(0.617)	(0.030)	(1.298)	(0.181)	(-3.357)	(-0.937)	(0.617)	(-2.892)	(0.013)**
TBOT	9.76	10.51	9.76	10.51	9.94	9.94	9.38	9.57	8.63	12.01	3.79
	(-0.190)	(0.381)	(-0.190)	(0.381)	(-0.043)	(-0.043)	(-0.490)	(-0.339)	(-1.126)	(1.426)	(0.925)
RTOP	9.38	9.38	10.07	11.21	12.13	11.44	9.61	12.59	6.18	8.01	14.97
	(-0.443)	(-0.443)	(0.047)	(0.804)	(1.363)	(0.947)	(-0.276)	(1.630)	(-3.318)	(-1.533)	(0.092)*
RBOT	4.07	9.30	11.05	10.85	13.57	9.11	12.21	10.27	12.02	7.56	34.04
	(-6.818)	(-0.546)	(0.758)	(0.623)	(2.366)	(-0.704)	(1.533)	(0.203)	(1.408)	(-2.099)	(0.000)**
DTOP	9.91	11.26	7.21	9.76	10.36	9.01	12.76	11.86	9.76	8.11	16.85
	(-0.078)	(1.030)	(-2.787)	(-0.209)	(0.305)	(-0.893)	(2.137)	(1.486)	(-0.209)	(-1.789)	(0.051)*
DBOT	8.70	11.06	9.88	10.62	10.91	11.65	9.29	9.59	9.88	8.41	6.78
	(-1.199)	(0.882)	(-0.103)	(0.524)	(0.764)	(1.341)	(-0.635)	(-0.365)	(-0.103)	(-1.495)	(0.660)

A.5 Kolmogorov-Smirnov Two Distributions Test

The table shows the Kolmogorov-Smirnov two distributions test of the equality of conditional and unconditional normlized 1-day return distributions for a sample of 65 stocks from OMX Swedish Large Caps from 1982 to 2006. The sample is divided into the full sample and 5 sub-periods. Conditional returns are defined as the daily return three days following the conclusion of an occurrence of one of ten technical patterns. P-values are with respect to the asymptotic distribution of the Kolmogorov-Smirnov test statistic. Added is also the condition of decreasing (Dec. Vol.) and increasing (Inc. Vol.) volume in our sample, as well as the difference (Diff.) between the two. Two asterisks indicates significance on the 5 percent level whereas one asterisk indicates significance on the 10 percent level.

Sample	**HS**	**IHS**	**BTOP**	**BBOT**	**TTOP**	**TBOT**	**RTOP**	**RBOT**	**DTOP**	**DBOT**
Full Sample	0.0026**	0.0070**	0.0127**	0.2182	0.0001**	0.2450	0.0562*	0.0142**	0.2943	0.3689
Dec. Vol.	0.0015**	0.0002**	0.0082**	0.2490	0.0114**	0.0172**	0.2079	0.1737	0.0238**	0.0018**
Inc. Vol.	0.0161**	0.0139**	0.0162**	0.0003**	0.1025	0.0042**	0.4178	0.0071**	0.0003**	0.0805*
Diff.	0.8069	0.3455	0.2709	0.1227	0.6893	0.4383	0.6947	0.6961	0.1255	0.4163
1982-1986	0.7774	0.6373	0.0770*	0.7017	0.6182	0.8487	0.3792	0.0040**	0.1941	0.5063
Dec. Vol.	0.1339	0.5936	0.0052**	0.2450	0.1804	0.8227	0.2743	0.8016	0.2132	0.7391
Inc. Vol.	0.1522	0.3641	0.6704	0.2399	0.0150**	0.3061	0.7436	0.0764*	0.2914	0.0577*
Diff.	0.7974	0.9988	0.0511*	0.1877	0.0581*	0.3005	0.2606	0.4536	0.9239	0.4232
1987-1991	0.7304	0.0926**	0.0889*	0.4564	0.4858	0.3607	0.9442	0.0032**	0.4372	0.7565
Dec. Vol.	0.3805	0.0120**	0.5053	0.4937	0.1105	0.8530	0.3100	0.2489	0.6697	0.2857
Inc. Vol.	0.0166**	0.5269	0.2860	0.2777	0.2764	0.4767	0.9314	0.3621	0.1375	0.9532
Diff.	0.6836	0.0481**	0.6317	0.4487	0.5741	0.9929	0.5305	0.5655	0.9379	0.7180
1992-1996	0.0502*	0.0077**	0.0923*	0.3071	0.1481	0.3643	0.3971	0.2389	0.5760	0.5387
Dec. Vol.	0.1098	0.4838	0.0533*	0.2391	0.0634*	0.0437**	0.1185	0.0807*	0.0879*	0.0673*
Inc. Vol.	0.0946*	0.2400	0.3064	0.0812*	0.0807*	0.0674*	0.2090	0.0401**	0.0758*	0.1345
Diff.	0.5383	0.9948	0.6740	0.9312	0.4498	0.0801*	0.6595	0.5246	0.0459**	0.2171
1997-2001	0.0010**	0.1458	0.0328**	0.2670	0.0481**	0.3199	0.1265	0.3094	0.0509**	0.0581**
Dec. Vol.	0.1325	0.3970	0.1023	0.3584	0.6856	0.0965*	0.3266	0.9010	0.5307	0.3296
Inc. Vol.	0.2408	0.3077	0.1493	0.5080	0.4519	0.3322	0.5488	0.1215	0.0240**	0.0180**
Diff.	0.9671	0.6684	0.0822*	0.6479	0.8241	0.9290	0.6253	0.3375	0.1018	0.1610
2002-2006	0.4042	0.1641	0.5125	0.1644	0.1703	0.1815	0.4255	0.0084**	0.1481	0.5527
Dec. Vol.	0.1972	0.0956*	0.0750**	0.0772*	0.0063**	0.0503*	0.6495	0.5441	0.0232**	0.0089**
Inc. Vol.	0.2129	0.0293**	0.5564	0.0999*	0.1101	0.0700*	0.5835	0.0179**	0.0452**	0.0628*
Diff.	0.8374	0.5644	0.9451	0.3236	0.1833	0.2694	0.9043	0.5793	0.5312	0.9576

B Appendix II - Graphical Representation

B.1 Head-and-Shoulders

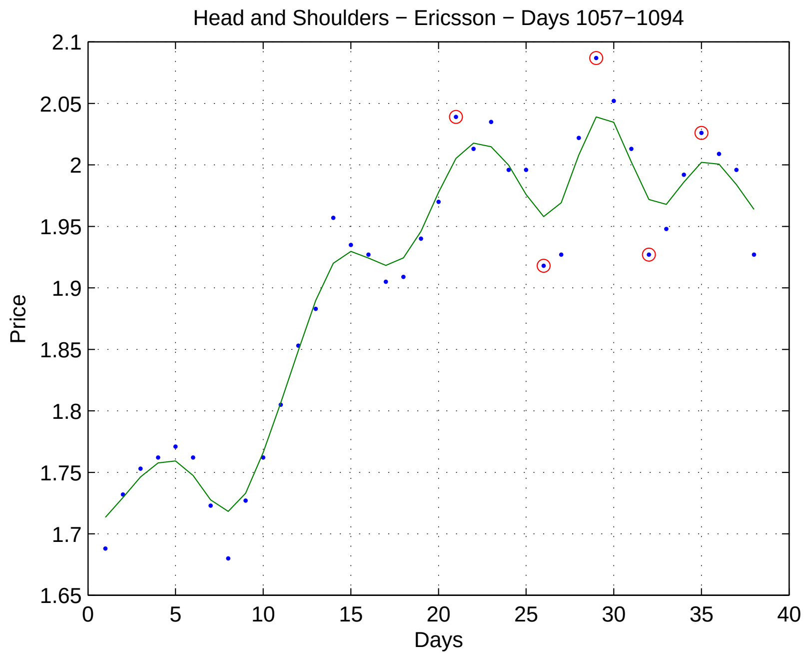

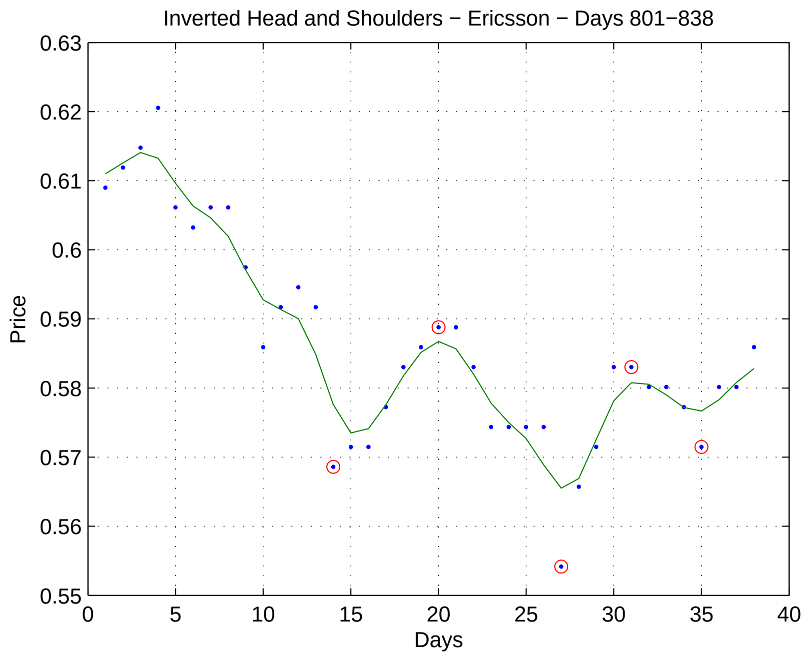

B.2 Broadening Tops and Bottoms

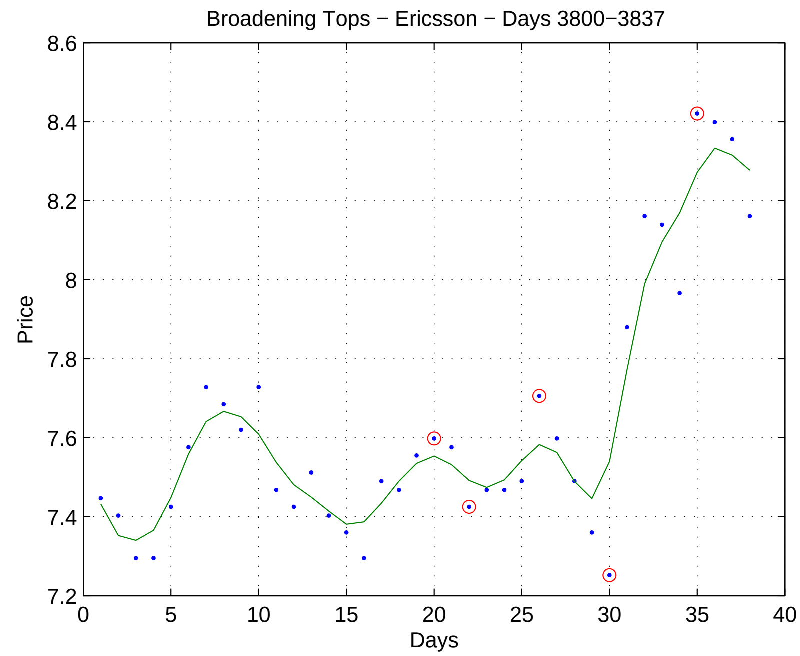

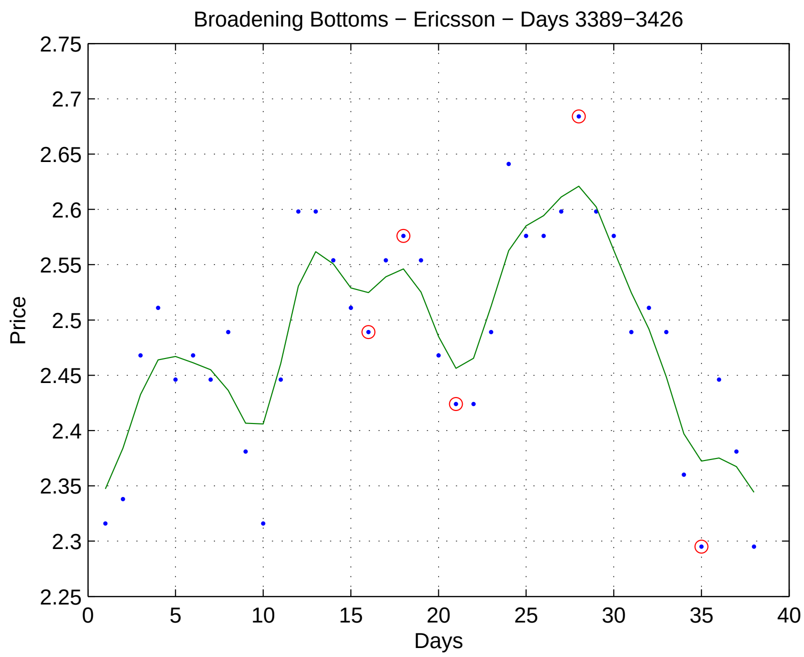

B.3 Triangle Tops and Bottoms

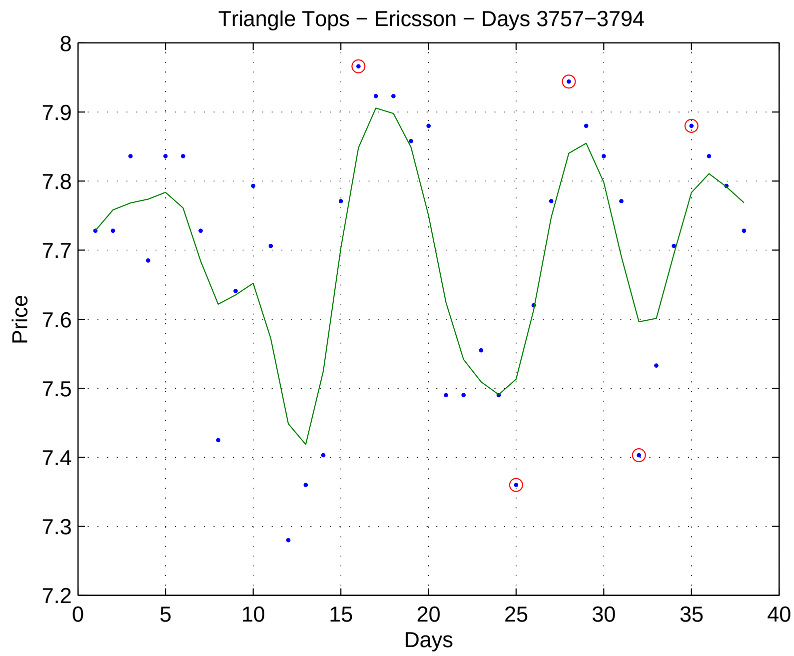

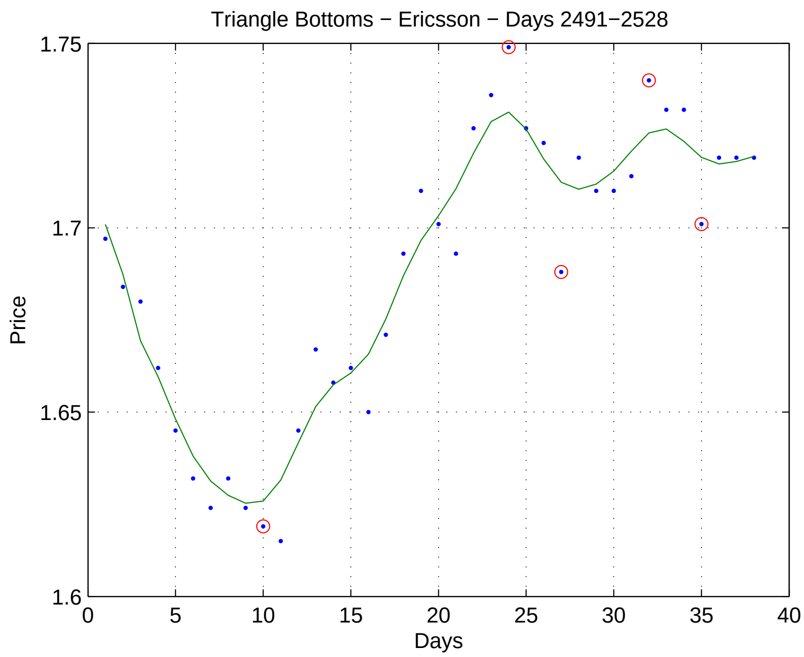

B.4 Rectangle Tops and Bottoms

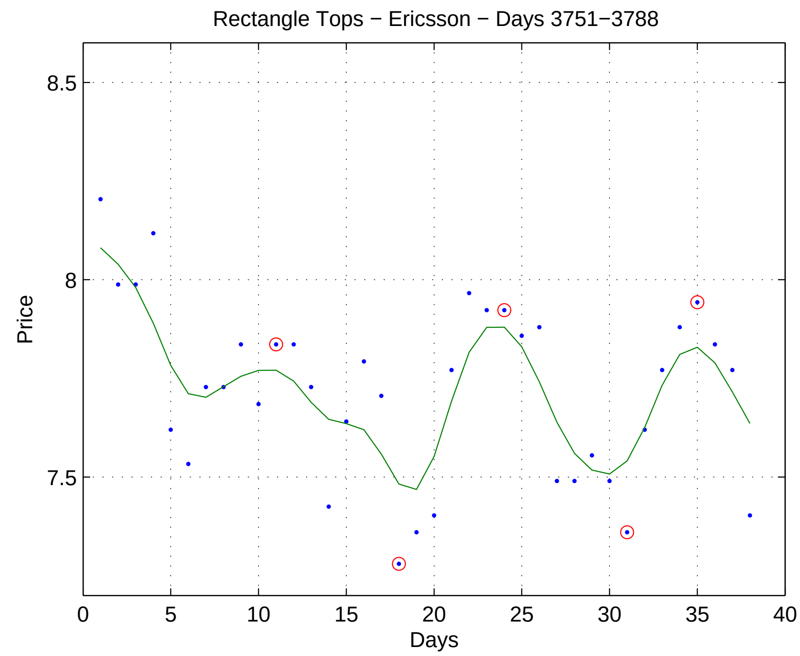

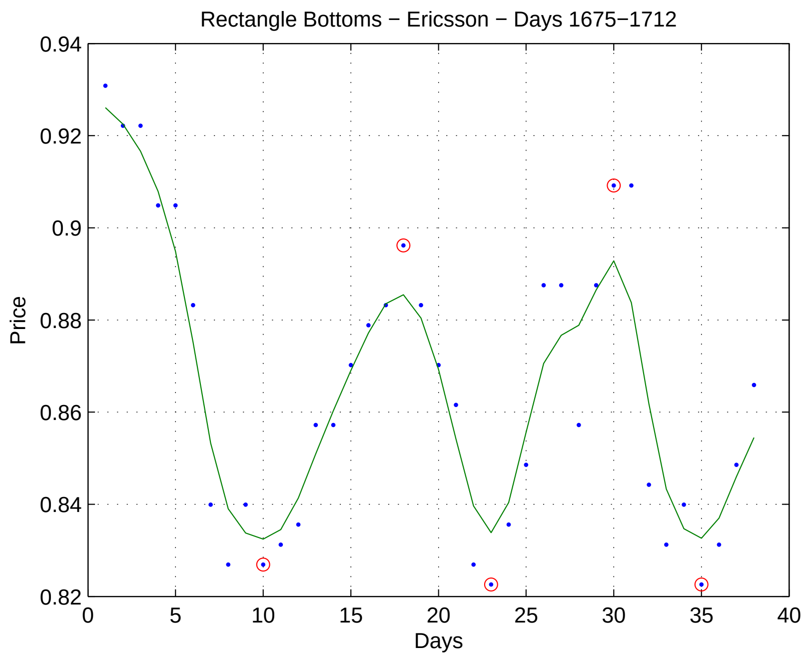

B.5 Double Tops and Bottoms